

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE
BEFORE THE BOARD OF PATENT APPEALS AND INTERFERENCES

Appellant: Lodha

Group Art Unit: 2616

Serial No. 10/620,668

Examiner: Su, Benjamin

Filed: July 16, 2003

Conf. No.: 9563

For: CUSTOMER-SPECIFIC TRAFFIC SHAPING

Mail Stop Appeal Brief-Patents

Commissioner for Patents

P.O. Box 1450

Alexandria, VA 22313-1450

APPEAL BRIEF UNDER 37 C.F.R. § 41.37(a)

This is an appeal to the Board of Patent Appeals and Interferences from the decision of the Examiner dated August 14, 2007, which finally rejected claims 1, 2, 4 – 7, and 15 – 20 in the above-identified application. The Office date of receipt of Appellants' Notice of Appeal was November 14, 2007. This Appeal Brief is hereby submitted pursuant to 37 C.F.R. § 41.37(a).

CERTIFICATE OF MAILING UNDER 37 C.F.R. 1.8

I hereby certify that this paper (along with any paper referred to as being attached or enclosed) is being facsimile transmitted to the Patent and Trademark Office on the date shown below.

Date: _____

Signed: _____

Typed Name: Mark A. Wilson

TABLE OF CONTENTS

I.	REAL PARTY IN INTEREST	2
II.	RELATED APPEALS AND INTERFERENCES	2
III.	STATUS OF CLAIMS	2
IV.	STATUS OF AMENDMENTS	2
V.	SUMMARY OF CLAIMED SUBJECT MATTER	2
VI.	GROUND OF REJECTION TO BE REVIEWED ON APPEAL	4
VII.	ARGUMENT	5
VIII.	CONCLUSION	10
IX.	CLAIMS APPENDIX	11
X.	EVIDENCE APPENDIX	16
XI.	RELATED PROCEEDINGS APPENDIX	17

I. REAL PARTY IN INTEREST

The real party in interest in this appeal is Alcatel-Lucent, 54, rue La Boetie, 75008, Paris, France.

II. RELATED APPEALS AND INTERFERENCES

To the best of Appellants' knowledge, there are no appeals or interferences related to the present appeal that will directly affect, be directly affected by, or have a bearing on the Board's decision in the instant appeal.

III. STATUS OF CLAIMS

Claims 1, 2, 4 – 10, and 14 – 20 are pending in the application and claims 1, 2, 4 – 7, and 15 – 19 were finally rejected in the Final Office Action mailed on August 12, 2007. In particular, claims 1, 2, 4 – 7, and 15 – 19 are rejected under 35 U.S.C. 102(e) as being anticipated by Bly et al. (U.S. Pat. Publ. No. 2004/0042399 A1, hereinafter Bly).

Claims 1, 2, 4 – 7, and 15 – 19 are the subject of this appeal. A copy of claims 1, 2, 4 – 10, and 14 – 20 as they stand on appeal is set forth in the Claims Appendix.

IV. STATUS OF AMENDMENTS

No claim amendments were filed subsequent to the Final Office Action mailed on August 14, 2007.

V. SUMMARY OF CLAIMED SUBJECT MATTER

This section of this Appeal Brief is set forth to comply with the requirements of 37 C.F.R. § 41.37(c)(1)(v) and is not intended to limit the scope of the claims in any way. Exemplary implementations of the limitations of independent claims 1 and 15 are described below.

Claim 1 relates to a method for forwarding packet-based traffic through a network node. Page 1, lines 13 – 15, page 2, line 11. Claim 1 recites receiving traffic type bandwidth limitations from a customer. Page 3, lines 26 – 28, page 5, lines 29 – 30, page 7, line 6, page 13, lines 6 – 7. Claim 1 also recites dedicating a group of queues in a

network node to the customer. Page 2, lines 12 – 13, page 5, lines 9 – 10, and Fig. 4B, 422. Claim 1 also recites translating the traffic type bandwidth limitations, which were received from the customer, to queue-specific bandwidth limitations. Page 7, lines 3 – 16. Claim 1 also recites performing queue-specific rate shaping on the customer's traffic according to the queue-specific bandwidth limitations respectively associated with the queues. Page 2, lines 13 – 14, page 5, lines 11 – 13, page 7, line 18. Claim 1 also recites performing group-specific rate shaping on the customer's traffic according to a group-specific bandwidth limitation associated with the group of queues. Page 2, lines 14 – 15, page 5, lines 13 – 14, page 7, line 19, page 4B, 430.

Claim 15 relates to a method for packet-based traffic forwarding. Page 1, lines 13 – 15. Claim 15 recites establishing a customer-specific bandwidth limitation for a customer. Page 3, lines 26 – 28, page 5, line 30. Claim 15 also recites receiving traffic-type-specific bandwidth limitations from the customer. Page 2, lines 12 – 13, page 5, lines 26 – 30, page 7, line 6, page 13, lines 6 – 7. Claim 15 also recites dedicating multiple traffic channels to the customer. Page 2, line 27, page 13, lines 9 – 10, Fig. 4B, 422. Claim 15 also recites associating the customer-specific bandwidth limitation to the traffic channels. Page 13, lines 16, Fig. 4B, 422. Claim 15 also recites associating the traffic-type-specific bandwidth limitations with the traffic channels. Page 7, lines 3 – 16. Claim 15 also recites performing traffic-type-specific rate shaping according to the traffic-type-specific bandwidth limitations respectively associated with the traffic channels. Page 2, lines 27 – 29, page 13, lines 17 – 18, Fig. 4B, 426. Claim 15 also recites performing customer-specific rate shaping according to the customer-specific bandwidth limitation associated with the traffic channels. Page 2, lines 29 – 30, page 13, lines 23 – 27, Fig. 4B, 430.

VI. GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL

Whether claims 1, 2, 4 – 7, and 15 – 19 are anticipated by Bly under 35 U.S.C. § 102(e).

VII. ARGUMENT

For the purposes of this appeal, claims 1 and 15 are argued together as a group.

Claim 1

Claim 1 particularly points out that traffic type bandwidth limitations are received from a customer and that it is the traffic type bandwidth limitations which were received from the customer that are translated to queue-specific bandwidth limitations.

Specifically, claim 1 recites:

“A method for forwarding packet-based traffic through a network node, comprising:
 receiving traffic type bandwidth limitations from a customer;
 dedicating a group of queues in a network node to the customer;
 translating *the traffic type bandwidth limitations*, which were *received from the customer*, to queue-specific bandwidth limitations;
 performing queue-specific rate shaping on the customer’s traffic according to the queue-specific bandwidth limitations respectively associated with the queues; and
 performing group-specific rate shaping on the customer’s traffic according to a group-specific bandwidth limitation associated with the group of queues.”
(emphasis added)

In the Final Office action, the “receiving” limitation was rejected in view of Bly paragraph [0025], line 1. Appellant asserts that Bly does not disclose “*receiving traffic type bandwidth limitations from a customer*” as recited in claim 1. At paragraph [0025], lines 1 – 3, Bly discloses “[T]he queues 44 – 47 can have shaping profiles, which include properties such as: priority, depth, latency, jitter, and rate.” Although Bly discloses that a queue can have a shaping profile such as rate, nowhere does Bly disclose that the shaping profile is received from the customer in the form of a traffic type bandwidth limitation. That is, while Bly discloses that a queue has a shaping profile, Bly does not disclose where the shaping profile comes from. Receiving traffic type bandwidth limitations from a customer and then translating the traffic type bandwidth limitations to queue-specific bandwidth limitations as recited in amended claim 1 allows the customer to control bandwidth usage without having to understand the concept of queuing within a service provider edge device. Because Bly does not disclose “*receiving traffic type bandwidth limitations from a customer*” as recited in claim 1, Appellant asserts that claim 1 is not anticipated by Bly.

In response to Appellant's response to the Final Office action, the Advisory action states:

"On page 2 of the remarks, the applicants submit that Bly does not disclose where the shaping profile comes from; therefore, Bly does not disclose 'receiving traffic type bandwidth limitations from a customer.' In reply, the examiner respectfully disagrees. Bly discloses that 'like traffic' can be defined as customer desired for a particular application. In order for the pay-per-view video traffic to be sent to a customer, the customer has to send traffic type bandwidth limitations to the node first, which reads on 'receiving traffic type bandwidth limitations from a customer.' The traffic type bandwidth limitations has no defined structure in the claim, therefore, it could be interpreted as broad as possible. In this case, the traffic type bandwidth limitations correspond to the pay-per-view request. Video traffic is given higher priority than non-real-time data, therefore, the bandwidth of the pay-per-view video traffic is higher than non-real-time data traffic (see paragraphs 24 – 25)." (emphasis added)

That is, the Examiner appears to argue that a pay-per-view request is equivalent to a traffic type bandwidth limitation from a customer. Appellant respectfully asserts that a pay-per-view request is not equivalent to a traffic type bandwidth limitation from a customer.

Claim 1 recites in part a traffic type "bandwidth limitation." By definition, a "bandwidth limitation" is a maximum limit on the amount of bandwidth that can be utilized/consumed. A pay-per-view request is a request to view content that requires an additional charge. A pay-per-view request says nothing about the maximum amount of bandwidth that the pay-per-view traffic can consume/utilize. A pay-per-view request is nothing more than an expression of a desire to buy content. Therefore, Appellant asserts that a pay-per-view request does not disclose "receiving traffic type bandwidth limitations from a customer" as recited in claim 1.

Additionally, giving video traffic a higher priority than non-real-time traffic says nothing about the maximum amount of bandwidth that the video traffic can utilize/consume. Giving video traffic a higher priority than non-real-time traffic does nothing more than establish a relative priority between video traffic and non-real-time traffic. Establishing a relative priority between different types of traffic does not disclose a bandwidth limitation for a certain traffic type. Because establishing a relative priority between different types of traffic does not disclose a bandwidth limitation for a certain traffic type, Appellant asserts that claim 1 is not anticipated by Bly.

Additionally, Appellant reiterates previous arguments that remain relevant to the issue at hand.

Firstly, paragraph [0024] of Bly discloses in full:

“[0024] The shaping engine 34 (see FIG. 4) en-queues incoming traffic 36 onto a selected one of the queues 44-47 based, for example, upon look-up information, which classifies the traffic. Streaming audio or video would be classified differently than e-mail, because streaming audio or video requires sufficient bandwidth to play without interruption. Therefore like-traffic, such as a stream or set of streams is placed in the same burst group 12, 14, or 16, in one embodiment. Within each burst group, further sub-classification can take place to determine on which one of the queues 44-47 the traffic 36 should be en-queued. “Like traffic” can be defined as desired for a particular application. It could be, for example, “all video traffic”, or it could be “all pay-per-view” video traffic, or it could be “all traffic for customer X”, or it could be “all email traffic.” It is a grouping of traffic with similar needs. Video, for example requires a fast rate, with low latency and jitter influences. Email on the other hand, can be handled on a “best efforts” basis; i.e. low-priority, without regard to latency and jitter.” (emphasis added)

As highlighted above, Bly discloses that like traffic can be defined as desired for a particular application. Bly then gives the example that like traffic could be identified as “all pay-per-view” video traffic, or it could be “all traffic for customer X”, or it could be “all email traffic.” Although Bly discloses different types of like traffic, Bly does not disclose who defines the like traffic. Specifically, Bly does not disclose that the customer defines like traffic types. It is more likely that the operator of the network (or the operator of the shaping engine in the case of Bly) is the entity that defines which traffic types are “like traffic” for traffic shaping/queuing purposes. Although Bly discloses that “like traffic” can be identified, Bly does not disclose that a customer is identifying which traffic types are “like traffic.”

Bly also discloses that the different types of traffic have different shaping profiles, i.e., fast rate, low latency, low jitter, best effort, low-priority. Although Bly discloses that the different traffic types have different shaping profiles, Bly does not disclose that the customer provides the specific shaping profiles to the network node. Appellant asserts that it is more likely the case that the customer provides the traffic to the network node, the network node identifies the traffic type, and then applies its own shaping profiles to the queues 44 – 47.

Secondly, the Advisory action states that:

“In order for the pay-per-view video traffic to be sent to a customer, the customer has to send traffic type bandwidth limitations to the node first, which reads on ‘receiving traffic type bandwidth limitations from a customer.’” (emphasis added)

Appellant respectfully disagrees that a customer has to send traffic type bandwidth limitations to a network node in order for pay-per-view video traffic to be sent to the customer. It is quite possible for a customer to request that pay-per-view video be supported by a network node without the customer providing a traffic type bandwidth limitation to the network node. For example, the customer can simply request the pay-per-view video content without specifying any particular bandwidth limitation at the network node. As is known in the field, a network node can be configured to recognize packet-based traffic as video traffic and handle the traffic to meet a certain pre-specified shaping profile. While the network node can recognize traffic as video traffic and can handle the traffic to meet a certain pre-specified shaping profile, the network node does not have to receive any type of bandwidth limitation from the customer as suggested in the Advisory action.

Additionally, paragraph [0025] of Bly discloses in full:

“[0025] The queues 44-47 can have shaping profiles, which include properties such as: priority, depth, latency, jitter, and rate. For example, video needs to always get through. A large amount of latency is not desirable for video, as any latency will cause the resulting picture to become jerky, and fall behind. The same is true of the rate at which video is sent. A constant, consistent stream should be used to supply the video information “just in time” for the next entry or element (e.g., packet or frame) of the picture on a TV or computer. Therefore, “video” traffic is properly classified so that it is managed appropriately. Because the video must always get through, it is given a “high” priority. Because video cannot be influenced/slowed-down with a large amount of latency, the depth of the queue is selected to be shallow. Therefore, little data can build up, waiting in the queue. With regard to rate, the video queue gets its own bandwidth end-to-end on a switch, and does not have to compete with any other queue for bandwidth. Queues for other classifications of traffic would similarly have appropriately chosen priorities, depths, latencies, jitter, and rates.” (emphasis added)

As highlighted above in paragraph [0025], Bly discloses that video traffic is given a high priority. Giving video traffic a high priority (e.g., a higher priority than non-real-time traffic) does nothing more than establish a relative priority between video traffic and another traffic type (e.g., non-real-time traffic). Establishing a relative priority between different types of traffic does not disclose a bandwidth limitation for a certain traffic type. Because establishing a relative priority between different traffic types does not disclose a

bandwidth limitation for a certain traffic type, Appellant asserts that claim 1 is not anticipated by Bly.

In view of the above-provided remarks, Appellant asserts that claim 1 is not anticipated by Bly.

Claims 2 and 4 – 7 are dependent on claim 1. Appellant asserts that these claims are allowable at least based on an allowable claim 1.

Independent Claim 15

Independent claim 15 includes similar limitations to claim 1. For example, claim 15 recites “receiving traffic-type-specific bandwidth limitations from a customer.” Because of the similarities between claim 1 and claim 15, Appellant asserts that the remarks provided above with reference to claim 1 apply also to claim 15. Appellant asserts that Bly does not disclose the above-identified limitations of amended claim 15.

Claims 16 – 20 are dependent on claim 15. Appellant asserts that these claims are allowable at least based on an allowable claim 15.

VIII. CONCLUSION

For the reasons stated above, claims 1, 2, 4 – 7, and 15 – 19 are patentable over the cited reference. Thus, the rejections of claims 1, 2, 4 – 7, and 15 – 19 should be withdrawn. Appellant respectfully request that the Board reverse the rejections of claims 1, 2, 4 – 7, and 15 – 19 under 35 U.S.C. § 102(c) and, since there are no remaining grounds of rejection to be overcome, direct the Examiner to enter a Notice of Allowance for claims 1, 2, 4 – 7, and 15 – 19.

At any time during the pendency of this application, please charge any fees required or credit any over payment to Deposit Account **50-3444** pursuant to 37 C.F.R. 1.25. Additionally, please charge any fees to Deposit Account **50-3444** under 37 C.F.R. 1.16, 1.17, 1.19, 1.20 and 1.21.

Respectfully submitted,

/mark a. wilson/

Mark A. Wilson
Reg. No. 43,994

Wilson & Ham
PMB: 348
2530 Berryessa Road
San Jose, CA 95132
Phone: (925) 249-1300
Fax: (925) 249-0111

Date: January 14, 2008

IX. CLAIMS APPENDIX

1. A method for forwarding packet-based traffic through a network node, comprising:
receiving traffic type bandwidth limitations from a customer;
dedicating a group of queues in a network node to the customer;
translating the traffic type bandwidth limitations, which were received from the customer, to queue-specific bandwidth limitations;
performing queue-specific rate shaping on the customer's traffic according to the queue-specific bandwidth limitations respectively associated with the queues; and
performing group-specific rate shaping on the customer's traffic according to a group-specific bandwidth limitation associated with the group of queues.
2. The method of claim 1, further comprising associating queues from said group of queues with different types of traffic that are to be received from the customer.
3. (canceled)
4. The method of claim 1, further comprising associating said group of queues with a group rate shaper that performs said group-specific rate shaping on the customer's traffic on an aggregate basis.
5. The method of claim 1, further comprising prioritizing the queues of said group of queues.
6. The method of claim 5, further comprising:
distributing said portion of excess unused bandwidth among the group of queues on a priority basis according to said prioritizing.
7. The method of claim 1, further comprising:

scheduling packets for forwarding from one or more of said queues in said group of queues, wherein bandwidth consumed by the packets from each of the queues is less than or equal to respective queue-specific bandwidth limitations for the queues;

identifying excess unused bandwidth when the consumed bandwidth is less than said group-specific bandwidth limitation; and

distributing a portion of the excess unused bandwidth to a first queue of the group of queues, wherein the sum of the consumed bandwidth and the portion of the excess unused bandwidth is less than or equal to a group-specific bandwidth limitation for the group.

8. A network node for forwarding packet-based traffic, comprising:

a plurality of queues;

a plurality of queue-specific rate shapers respectively associated with the plurality of queues;

a plurality of group-specific rate shapers configured to be associated with groups of the plurality of queues;

a group establishment module configured to dedicate a group of said queues to a customer and to associate one of said group-specific rate shapers with said group of queues that is dedicated to said customer; and

a scheduler configured to:

schedule, in a first round, packets enqueued in the plurality of queues according to the respective plurality of queue-specific rate shapers and the respective group-specific rate shapers; and

schedule, in a second round, packets enqueued in the plurality of queues according to the respective group-specific rate shapers;

wherein said scheduler is further configured to:

schedule, in subrounds of the first round, packets enqueued in the plurality of queues according to a priority respectively associated with each of the queues and schedule, in subrounds of the second round, packets enqueued in the plurality of queues according to the priority respectively associated with each of the queues;

wherein the scheduler comprises an individual queue enablement vector for each queue, a group enablement vector for the group of queues, and a result vector for each queue;

wherein the individual queue enablement vector indicates which queues are enabled, with a queue being enabled if the queue has not consumed its allocated queue-specific bandwidth;

wherein the group enablement vector indicates whether the group is enabled with the group being enabled as long as all of the allocated group-specific bandwidth has not been consumed; and

wherein the result vector indicates which queues are enabled for sending packets, wherein in the first round a result vector for a queue indicates a queue is enabled only when both the individual queue enablement vector and the group vector indicate that the queue is enabled and in the second round a result vector for a queue indicates a queue is enabled as long as the group vector indicates that the group is enabled.

9. The device of claim 8, further comprising:

a scheduler, coupled to the plurality of queue-specific rate shapers and the plurality of group-specific rate shapers, configured to schedule packets enqueued in the plurality of queues according to the respective plurality of queue-specific rate shapers, wherein the queue-specific rate shaper respectively associated with each queue is associated with a priority, and wherein the scheduler schedules according to the associated priority.

10. The device of claim 9, wherein said scheduler is further configured to:

scheduling packets for forwarding from a first one or more queues of said plurality of queues, wherein bandwidth consumed by the packets from each of the first one or more queues is less than or equal to respective queue-specific bandwidth limitations for the first one or more queues;

identifying excess unused bandwidth when the consumed bandwidth is less than a group-specific bandwidth limitation, wherein a sum of the consumed bandwidth and the excess unused bandwidth approximately equals the group-specific bandwidth limitation; and

scheduling packets for forwarding from a second one or more queues of said plurality of queues using the excess unused bandwidth.

11. (canceled)

12. (canceled)

13. (canceled)

14. The device of claim 8, further comprising:

a plurality of pipes, wherein each pipe is associated with a group-specific rate shaper, and wherein each pipe of said plurality of pipes includes:

multiple traffic channels comprising one or more queues of the plurality of queues, wherein each traffic channel is associated with a queue-specific rate shaper.

15. A method for packet-based traffic forwarding, comprising:

establishing a customer-specific bandwidth limitation for a customer;
receiving traffic-type-specific bandwidth limitations from the customer;
dedicating multiple traffic channels to the customer;
associating the customer-specific bandwidth limitation to the traffic channels;
associating the traffic-type-specific bandwidth limitations with the traffic channels;
performing traffic-type-specific rate shaping according to the traffic-type-specific bandwidth limitations respectively associated with the traffic channels; and
performing customer-specific rate shaping according to the customer-specific bandwidth limitation associated with the traffic channels.

16. The method of claim 15, further comprising:

prioritizing the traffic channels relative to one another.

17. The method of claim 16, wherein said performing traffic-type-specific rate shaping consumes less bandwidth than said customer-specific bandwidth limitation, said method further comprising:

identifying excess unused bandwidth following the traffic-type-specific rate shaping;
and

distributing the excess unused bandwidth to a subset of the traffic channels in priority order according to said prioritizing.

18. The method of claim 15, further comprising:

associating a traffic type with each traffic channel.

19. The method of claim 18, further comprising:

adjusting the traffic-type-specific rate shaping according to traffic type-specific rate shaping customer preferences.

20. The method of claim 15, further comprising:

associating respective traffic-type-specific bandwidth limitations with each traffic channel, wherein the sum of the respective traffic-type-specific bandwidth limitations is less than or equal to the customer-specific bandwidth limitation.

X. EVIDENCE APPENDIX

There is no evidence submitted with this Appeal Brief.

XI. RELATED PROCEEDINGS APPENDIX

To the best of Appellant's knowledge, there are no appeals or interferences related to the present appeal that will directly affect, be directly affected by, or have a bearing on the Board's decision in the instant appeal.